Draft Minutes
Advanced Scientific Computing Advisory Committee Meeting
Oct. 28-29, 2008, Hilton Hotel, Gaithersburg, Md.

ASCAC members present:

F. Ronald Bailey
Marsha J. Berger
Jacqueline H. Chen
Roscoe C. Giles, Chair
James J. Hack
Thomas A. Manteuffel
John W. Negele
Horst D. Simon
Rick L. Stevens
Victoria A. White
Thomas Zacharia

ASCAC members absent:

David J. Galas
Anthony J. G. Hey
Larry L. Smarr
Robert G. Voigt

Also participating:

Melea F. Baker, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Christine A. Chalk, ASCAC Designated Federal Officer
Lali Chatterjee, Office of Advanced Scientific Computing Research, Office of Science, USDOE
S. Scott Collis, Technical Manager, Sandia National Laboratories
Vincent Dattoria, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Paul Fischer, Mathematics and Computer Science Division, Argonne National Laboratory
Omar Ghattas, Departments of Geological Sciences and Mechanical Engineering, University of Texas
Garth Gibson, Carnegie-Mellon University and Panasas, Inc.
Barbara Helland, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Bruce Hendrickson, Senior Manager for Computer Science and Mathematics, Sandia National Laboratories
Alexandra Landsberg, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Tom Lange, Director of Modeling and Simulation, Proctor & Gamble
Frederick M. O'Hara, Jr., ASCAC Recording Secretary
Walter M. Polansky, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Michael R. Strayer, Associate Director, Office of Advanced Scientific Computing Research, Office of Science, USDOE

About 50 others were in attendance in the course of the two-day meeting.

Tuesday, October 28, 2008
Morning Session

Before the meeting, new members of the Committee were sworn in and given a briefing on ethics. Chairman **Roscoe Giles** called the meeting to order at 9:09 a.m. He asked the Committee members to introduce themselves and reviewed the goals of the Committee.

Michael Strayer welcomed the new members and was asked to report on the status of the Office.

The Federal Government is operating under a continuing resolution (CR) until March 2009, which makes budgeting and planning hard. Both Senate and House have an Energy and Water Bill that has passed the respective Committee. Those bills restrict the Office of Advanced Scientific Computing (ASCR) to roughly the FY08 spending levels. After the bills are passed by the full Senate and House, the two committees go to conference and then back to both houses for votes before going to the President to become law. All of this needs to happen before January 1, or it starts all over again. The new FY10 budget request is to be delivered in January. There is a probability that the CR will be extended to a full year.

The House report recommended $378,820,000 for ASCR, an increase of $10,000,000 over the budget request and $27,647,000 over the current fiscal year. The increase includes $5,000,000 above the budget request to expand its Innovative and Novel Computational Impact on Theory and Experiment (INCITE) activities. A further $5,000,000 is provided to enhance advanced scientific computing research activities relevant to two of the six integrated research and development areas identified in the request; $5,000,000 is provided for Advanced Mathematics for Optimization of Complex Systems, Control Theory, and Risk Assessment; and $2,969,000 is provided for Carbon Dioxide Capture and Storage.

The Senate report provides a little less money for ASCR, $368,820,000, the same as the budget request. The Senate Committee was concerned that the Department has limited cooperation between the National Nuclear Security Administration (NNSA) and DOE laboratories in supporting the advanced computing architecture and algorithm development. The Department has stepped forward to make it known that the national laboratories have a vigorous collaboration with NNSA.

The Applied Mathematics budget was to go from $36.885 million to $46.164 million but was held to $36.885 million by the CR. The Computational Partnerships Program is down a bit. The Next-Generation Networking for Science is being held to its FY08 level instead of being built out. High Performance Production Computing and the Leadership Computing Facilities (LCFs) are down a bit but not dramatically. Research and Evaluation Prototypes have more funds available in FY09 under the CR than was planned. These values bring the total ASCR FY09 budget to $351.173 million.

A year-long CR will delay the Joint Applied Mathematics-Computer Science Institute, slow down the solicitation for mathematics of large datasets, delay direct support for science-application “leading-edge developers,” affect the partnership with the Office of Biological and Environmental Research (BER) to improve climate models, and delay basic research in Cyber Security for Open Science. However, the CR fully funds the Department's commitment to the Defense Advanced Research Projects Agency's High Productivity Computing Systems.

ASCR continues to be understaffed, even as funding increases; the organization chart still has a lot of empty positions. The organization will be based on teams.

The Multiscale Mathematics and Optimization for Complex Systems call funded 2 proposals in FY08, and 19 proposals are "under consideration" pending the resolution of the CR. An Applied Mathematics meeting was held in Chicago on Oct 7-9, 2008, with parallel tracks on (1) joint mathematics/computer science institutes and (2) high-risk/high-payoff technologies.

The National Science Foundation (NSF) has undergone personnel changes, and joint programs are being discussed with the NSF Division of Mathematical Sciences.

The Applied Mathematics 2008 principal-investigator (PI) meeting was held at Argonne National Laboratory Oct 15-17, 2008, with more than 140 researchers in attendance.

In Computer Science Research, the Software Development Tools for Improved Ease-of-Use of Petascale Systems call produced 97 proposals, representing 34 projects. The review was held August 26-27, 2008.

Fred Johnson has developed a plan to strengthen ASCR's collaboration with NNSA.

The DOE Computational Science Graduate Fellowship Program is a hallmark of this Office to educate the next generation of computational scientists. The 2006 program review said that "This relatively small, but incredibly effective program has succeeded in the critical area of advanced scientific computing by operating a program that attracts and selects students through a competitive process that results in an enhanced graduate education in this important field…The success of this program is clearly evident." That review recommended that DOE should seek funding to double the size of the current program and should work closely with the contractor to be sure that the current excellent management approach is maintained. The notion of starting follow-on programs, such as named postdoctoral fellowships or young-investigator awards in scientific computing is worthwhile and should be considered. But it is important that the existing graduate student efforts not be diluted. The Office will continue a strong emphasis on this program. Currently, there are 68 fellows; 18 new fellows started in September 2008. The next annual conference for fellows will be July 14–16, 2009, in Washington, D.C.

The Research Alliance in Math and Science (RAMS) was designed to provide collaborative research experiences among faculty and students at colleges or universities and DOE national laboratory researchers. Each student is assigned to a research mentor.

At the petascale, a good job has been done in standing up and upgrading facilities. A very good job has not been done on the support software. The full buildout of libraries and operating systems has not been done because of a lack of money and manpower. Open-source software is being developed for the petascale machines. A multiagency approach will go forward in collaboration with European and Asian partners to produce software that will run very efficiently on these new machines. High efficiency and productivity is demanded, given the size and cost of these machines. Workshops will be held on open software and on best practices.

The INCITE Program was the idea of DOE's Under Secretary for Science. It was initiated at the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL) in 2004. Office of Science (SC) computing resources are being provided to a small number of computationally intensive research projects of large scale that can make high-impact scientific advances through the use of a large allocation of computer time and data storage. INCITE is open to national and international researchers, including industry; 80% of the LCFs will be allocated through INCITE. For FY09, 79 proposals and 21 renewal proposals were received. More than 2 billion processor hours were requested from new and renewed proposals, and more than 600 million processor hours are available for FY09 awards, including renewals.

The Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR) programs operate as a separate division in ASCR. The current set-asides for SBIR and STTR are 2.5% and 0.3%, respectively, on all DOE R&D programs except NNSA. A new SBIR/STTR request for proposals (RFP) was issued September 17, 2008, and will close November 20. The selection of 350 SBIR Phase I and 25 STTR Phase I awards (at $100,000 each) in April is anticipated. ASCR-related topic areas can be developed in a small-business environment and built up by DOE.

The ASCR PI Meeting will be held in Austin, Texas, on November 17, 2008.

Giles asked what the opportunities would be for public engagement in the strategic-planning efforts. Strayer replied that the communities are being asked to develop plans that will then be integrated. The final plan will come to ASCAC for review.

Berger noted that a utilization rate of 90% had been cited for one of the LCFs and suggested that a better metric might be the time to solution for an algorithm. Strayer agreed. Cutting forefront science is pushing high-efficiency algorithms.

Negele asked if there were movement on postdoctoral fellowships. Strayer replied that there was. There is now a slate of fellowships, but it is difficult to fund under the continuing resolution. This slate of fellowships will be built out under future budgets.

White asked about ASCR's role in the leadership of Scientific Discovery Through Advanced Computing (SciDAC). Strayer responded that he was also the acting SciDAC program manager. About $50 million is invested in SciDAC. The program was recompeted, and is now focusing on new instruments. It is very successful. It contributes significantly to breakthrough science. It is in the forefront of developing the extreme scale. The NSF wants to be a partner in future solicitations, which will increase investment in this program. Investments also need to be made in data.

Simon asked what the plans were for exascale computing. Strayer replied that town meetings were held; the report of those town meetings has led to a series of workshops (which will start the week after this meeting) to find out the large societal problems whose solutions can be contributed to by exascale computing. Discussions are also being held by program managers at all levels, including a joint ASCR−NNSA workshop. Advances in hardware and software will be needed, and the ability to do ensemble computing [up to one million central processing units (cpus)] will need to be developed.

A break was declared at 10:20 a.m. The meeting was reconvened at 10:39 a.m. Giles began a discussion of the "balance" report. There was a long-standing charge to assess the balance between investments in facilities and research. A previous consideration of the Balance Subcommittee's report remanded the draft to the Subcommittee for changes. A revised draft was circulated to the Committee two weeks before this meeting. Additional comments have been submitted since then.

Reporting for the Subcommittee, **Ronald Bailey** noted that ASCR has a large investment in hardware. Investments in research in computer science should be equal to those in hardware. To get into the game, one needs the hardware. To win the game, one has to have the software. The second balance considered is between short- and long-term research. One needs to produce seed corn for later crops. If one looks at INCITE and other extreme computing, those projects are making breakthroughs that could not be made otherwise. Resources are needed to do this extreme computing, which produces a lot of science. Congress recognizes that one has to have a lot of computing resources to do the science (like producing fusion power).

Manteuffel noted that having petascale and exascale computers is a necessary but not sufficient requirement for getting science done. One cannot neglect algorithms, computing science research, models, etc. In his estimation, the report was well done.

Berger agreed with Manteuffel. The editing took care of her previous concerns. However, in basic application research, the discussions focus on the LCFs and should be broader. A lot of the LCF run times are for the broader DOE community, and the report could more specifically describe the portfolio and its holders. Bailey replied that there are new problems (e.g., multiscale problems), and new mathematical capabilities are needed. Giles noted that the report attempts to address the need for new mathematics for new applications.

Negele said that he appreciated the effort that was put into this report. Reading it as an outsider, it seemed to him to back off from its conclusions. The investment in leadership-class computers has been stunning and must continue. Pushing new architectures forces all the gains.

Chen said that extreme computing is for heroic runs. Parametric studies at scale should be added, blurring the distinction between exploitative and extreme computing. Strayer pointed out that Congress mandated that DOE computers be used for all the government.

Stevens agreed with Negele's comments and went on to say that ASCR has a broad mandate to provide computing leadership for the whole nation. There is a huge sea change occurring that is not recognized by the report: DOE is in the driver's seat. A strategy should be charted out to maintain the U.S. leadership. Three things need to be emphasized: (1) maintaining leadership across the board, (2) investment in science through investments in the computational domains, and (3) a DOE roadmap for creating a next generation of hardware during the next decade.

Simon stated that the report distinguishes between exploitation and extreme computing. That distinction does not exist. There is a fuzzy barrier between the two, and the distinction is not necessary. Also, at the exaflop, the United States cannot overcome the challenges alone; it will need to have international partners.

Zacharia agreed with Negele. It is important that the mission of SC and of science in general be debated. The LCFs have been cited as the most important science facilities for the nation. If one looks at the ASCR budget back to 2000, one finds that the ratio for research to facilities was $42 million to $78 million; today it is $151 million to $217 million. Research has grown by a factor of 3.54, and facilities have grown by a factor of 3. However, only a portion of SciDAC research is captured in these numbers. The balance ebbs and flows, but a good level of balance has been maintained. SciDAC's proposed hardware component did not get funded. The report should call for more mathematics and computer science but also should call for support of large-scale science.

White was struck by (1) the lack of an economic model of investments and a picture of how one might go forward and (2) a model of what will be required to do competitive science in the future. The report does not address the complex nature of the world and the balance between competitive and collaborative work.

Bailey said that the Subcommittee did not look on this report as a plan. ASCR should put together a strategic plan and ask ASCAC to review it. The report suggests that a model be constructed to guide the determination of the balance of resources. The Subcommittee did not get into the in-house/outside debate. Also, it did not want to write a report that asked for more money. The report *should* identify ASCR strategic objectives. Perhaps ASCR should be the leader, and the other DOE offices should support what benefits them.

Manteuffel said that the Committee has heard what is going to happen in fusion research, climate change, etc. and also about exascale computing. One needs to do the homework on how to build such a machine. Stevens pointed out that building a machine is a long-term process. One has to pursue *all* of these things and do it in a way that stimulates the next generation of researchers.

Giles suggested reshaping the question, reshaping the report to answer that question, and revisiting this topic on the second day of the meeting. Chalk announced revisions to the agenda to accommodate that discussion.

Garth Gibson was asked to discuss the Petascale Data Storage Institute, part of the SciDAC Program.

The high-level goal is to map out the needs for storage at the exascale. A growth of 100% per year is being seen in the storage required to keep up with growth in processing power, For disks, this means a 20% per year increase in bandwidth, one-fifth of the overall storage growth.

The SciDAC Petascale Data Storage Institute draws on expertise and experience at a number of institutions. The Institute's efforts are divided into three primary areas: outreach and leadership, data collection and dissemination, and mechanism innovation to solve some of the problems. It runs three forums each year on supercomputing, file systems and I/O, and file storage technology.

A persistent problem is the investment in scalability. Innovative file structures are constantly needed. The next-generation network file structure (NFS) is parallel (pNFS), but that puts everything through a network address. What is needed is a map to the network. Also needed are open source and competitive offerings. Those are coming along; a number of companies are involved; and a standard will be in place in a few months.

Los Alamos has kept root-cause fault-data logs on 22 clusters and 5000 nodes The data cover 9 years and continue. This effort will allow modeling failure in large-scale systems. The hope is that the failure rate per chip can be held constant. With these data, a simple model was constructed. Both the number of chips and the number of cores per chip will increase. The mean-time interrupt will drop because of the size of the system. An image can be made periodically, and processing can roll back to that image when a fault occurs. The application utilization will crash because checkpoints will have to be performed more and more frequently. Storage bandwidth will have to increase 130% per year because of the 25% increase in chips. The best solution is to compress the checkpoints and represent the state in denser terms to counter this effect. Without compression, a dedicated checkpoint device will be needed. Or, one could nail this problem with process-pairs duplication.

Storage suffers failures, too, so data on disk failures were also collected.

Fixing a drive means replacing all its content. But disks are getting bigger, and recovery per failure is therefore increasing. Maybe soon, hundreds of concurrent recoveries will be needed. The normal case must be designed for many failures, a huge challenge.

One can defer the problem by making failed disk repair a parallel application through file replication; and, more recently, object RAID [redundant array of inexpensive disks] can scale repair.

In summary, data on data storage are being gathered continually, file systems are being nurtured that match high-performance computing (HPC) scale and requirements; checkpoint specializations are being provided; failure is being understood as the normal case; correctness at increasing scale is being dealt with through testing and formal verification; and HPC vs cloud storage architecture is being looked at.

Stevens asked what an agency can do to build a more robust community for cloud solutions. Gibson replied, interest the academic community. The cloud world builds the computer it needs for capacity computing. Stevens asked what to do about research. Gibson answered, seduce people working on clouds, and they will work on your applications.

Negele asked if there were any new storage technologies Gibson replied that flash has made the transition and that magnetoresistive random-access memory (MRAM) would make a big contribution.

Simon asked if there were any research on data. Gibson answered that the dynamic range in file size has increased. If indexing and searching are increased, metadata will not scale up.

The floor was opened for public comment. There being none, a break for lunch was taken at 12:03 p.m.

Tuesday, October 28, 2008
Afternoon Session

The meeting was called back into session at 1:29 p.m. with the introduction of **Tom Lange** to present the perspectives of an INCITE participant. He described the business products and history of his company, Procter & Gamble.

Products must perform as expected (advertised) when used. They are made of materials that are strong but soft, must stretch but not break, must breathe but contain, and must break but not tear. Liquid mixtures may not separate, must dispense easily, and must stay where applied. Packages must be strong but light and must never leak but open easily.

Modeling and simulation has transformed the industries of defense, aviation, automotion, and consumer goods from atoms to the enterprise, and HPC is behind each of those processes. Modeling and simulation consist of computational chemistry, computer-aided engineering, empiricism, and optimization. They are used for supply-chain analysis, planning and scheduling analysis, plant throughput analysis, reliability engineering (10,000 bottle analyses last year), consumer-response modeling, optimization, process-reliability analysis, and solid mechanics (studying rigid-body kinematics with finite-element analysis).

Problems analyzed with modeling and simulation include checking bottles' squeeze performance, optimizing the bottle weight, packing performance, and racetrack performance. Computational fluid dynamics (CFD) is used to analyze free surface flow, contained turbulent flow, multi-phase flows, creeping and low Reynold's number flows, non-newtonian and visco-elastic material properties, flow in porous media, mixing of viscous and nonviscous liquids, air entrapment during bottle filling, and fluid structure interaction.

Computational chemistry disciplines include atomistic and quantum methods, molecular dynamics, meso-scale methods, quantitative structure-activity relationships/quantitative structure-property relationships (used in toxicology), quantum chemistry, and stability of complex fluids (toothpaste is really bad if one cannot squeeze it out of the tube or it runs off the toothbrush).

Multi-scale modeling is used to analyze surfactant lather for shampoos. The need is to predict equilibrium surface structure, properties, and behavior to understand micellization, interfacial effects, and "soil" removal.

Computing hardware performance is increasing rapidly. P&G's 1st, 2nd, and 3rd Generation machines are a generation behind DOE leadership-class machines. Computing costs are going down fast. With all that power, one can pursue realism and replace full-scale and -speed tests by solving bigger, more complex problems; doing parametric studies vs. point estimates; and reaching more analysts. (Of the 10,000 chemists and engineers at P&G, 2% are in modeling and simulation.)

The challenges include application software, data management, and education.

Simon asked if the application-software challenge should be outsourced. Lange answered that Procter & Gamble would say, "Let our people make consumers happy." If one can buy software rather than writing programs, do it. But, one can have commercial software available but not be able to afford to purchase it, so one writes it oneself, solving the same problems again.

Giles asked what Procter & Gamble's experience in application software was in INCITE. Lange replied that it was a very good experience. The partnership with an academic and a national laboratory worked extremely well for Procter & Gamble.

Manteuffel asked if they found the software developed at the national laboratories useful. Lange said that they would prefer to buy a commercial product rather than building an application. He missed the Red Hat connection. Some open-source programs are useful if they are going to be used repetitively.

Stevens asked how Lange tried to value these services to the corporation. Lange said that some tests done now are going to be replaced with simulations. Beakers are going to be replaced with chips. Also some product solutions that would never have been thought of will be found, and the time to market will be reduced by half. These simulations will save the corporation cash. Examples and experiences will be used to back up these claims. Stevens asked Lange how he dealt with his experimental brethren. Lange said that, in the old days, modelers got called on when experimentalists did not understand something. Today, modelers can be helpful in understanding phenomenology and by being predictive. If personnel can be retrained to do simulations rather than bottle-drop tests, their demise is not being caused by modeling and simulations.

Vincent Dattoria was asked to report on ASCR's Facilities Division.

DOE is doing leadership computing because the DOE High-End Computing Revitalization Act of 2004 required that DOE carry out a program of research and development (including development of software and hardware) to advance high-end computing systems and develop, deploy high-end computing systems for advanced scientific and engineering applications, and reduce risk by platform and geographic diversity. The ASCR facilities strategy entails providing the tools (high-end computing), investing in the future (research and evaluation prototypes), and linking it all together with the Energy Sciences Network (ESnet).

NERSC is located at LBNL. It has the Cray XT-4 Franklin (102 Tflop/s, 9,660 nodes, and 19,320 cores), the only machine that has a Program Assessment Rating Tool (PART) metric (40% of its computing time is to be used for large projects). Franklin is being upgraded to 350 Tflop/s and 38,640 cores in November. There is also a NERSC-6 Project, the RFP for which was issued in September 2008. NERSC also has an IBM Power 5, a Linux Opteron Cluster, and a Parallel Distributed Systems Facility Linux Cluster. 90% of NERSC resources go to SC users, and 10% to INCITE. NERSC is widening scientific discovery in validating climate models, simulating a low-swirl burner fueled with hydrogen, astrophysical plasmas, and nanoscience calculations and scalable algorithms. Many of the codes that start out at NERSC become candidates for INCITE use.

The Argonne LCF has a 111-peak-teraflop IBM Blue Gene/P with 8,192 quad-core compute nodes and 16 TB of memory. An upgrade to the machine was accepted in March 2008 and is in transition to operations, which will result in a 556-peak-teraflop machine with 40,960 quad-core compute nodes and 80 TB of memory. More than 400 million hours will be allocated to INCITE on the Argonne LCF for research on bubble formation, Parkinson's disease, fission-reactor design, and cardiac simulation.

The LCF at Oak Ridge has a 263-teraflop Cray XT4 with 7,832 quad-core, 2.1-GHz AMD-Opteron compute nodes and 62 terabytes of aggregate memory. It also has an 18.5-teraflop Cray X1E with 1,024 multi-streaming vector processors. Delivery of a 1-petaflop Cray is expected to be accepted by December 2008. Its scientific discovery work includes electron pairing in high-temperature superconducting (HTc) cuprates, shining the light on dark matter, stabilizing a lifted flame, taming turbulent heat loss in fusion reactors, and how a pulsar gets its spin. The facility uses a liaison approach, where a scientist works with staff to optimize the code.

Steve Cotter is the new head of ESnet. There is close collaboration between ASCR network research and ESnet technology. The ESnet 4 build-out is nearing completion. The Science Data Network (SDN) has 16 of 17 nodes deployed The metropolitan area networks (MANs) will be upgraded before the Large Hadron Collider (LHC) is brought back on line. The first LHC particle beams employed the largest computer grid in the world. There is now extensive international connectivity and collaboration.

On August 19-20 operational assessments were performed for the Argonne LCF, ESnet, NERSC, Oak Ridge LCF, and Molecular Science Computing Facility to review facility performance and the plans for the operational phase. All the reviews of the facilities were positive.

The 2008 "Best Practices" Workshop focused on Risk Management Techniques and Practices for High-Performance Computing Centers. It had about 70 participants, was jointly sponsored by SC and NNSA, was hosted by Lawrence Livermore National Laboratory (LLNL), and was held September 17-18 in San Francisco. Its report is due in November 2008.

Negele asked if there were any plans for next-generation LCFs. Helland replied that the Office had received mission statements for upgrades for both LCFs.

Giles asked how badly the facilities were hit by the continuing resolution. Dattoria replied that spending plans have been projected for all facilities, and they are living within those constraints.

Paul Fischer was asked to discuss the simulation of advanced nuclear reactors on LCF machines at Argonne National Laboratory (ANL).

Current light-water-reactor designs have been tweaked already. By burning minor actinides, fast reactors offer the potential of a 99% reduction in geological-repository requirements. DOE's Office of Nuclear Energy (NE) has recently embarked on an ambitious simulation program for reactor modeling, reprocessing, seismic analysis, etc. ANL is focusing on neutronics and thermal hydraulics.

The thermal hydronics modeling approach is a multiscale simulation hierarchy involving experiments, direct numerical simulation (DNS) of turbulence, large eddy simulation (LES), Reynolds-averaged Navier-Stokes (RANS), and subchannel or lumped-parameter models. Two problem areas identified were (1) mixing and pressure drop in fuel rod bundles, and (2) thermal mixing in the upper plenum.

Thermal mixing in the upper plenum influences longevity of mechanical structures, places reactor design constraints on outlet temperature differences, and is not well-understood. ANL is investing $1 million in a detailed experiment; Blue Gene/P simulations are supported through INCITE.

DNS of a simple pin model, RANS of up to 217 pins, and subchannel analysis are being used to analyze the fuel bundle subassembly. Interchannel cross-flow is the principal cross-assembly energy transport mechanism. The research team is in a position to solve this numerically so the assembly never needs to be built. Turbulent flow in a reactor sub-assembly with 37 wire-wrapped fuel pins has been simulated, and LES and RANS simulations gave comparable results in a 7-pin case.

On the applied-mathematics side, ANL is trying to enable advanced scientific simulation at the petascale and beyond by developing state-of-the-art algorithms and discretizations that are high-order, stable, and scalable. The physics focus is on fluid mechanics, heat transfer magnetohydrodynamics (MHD), and electromagnetics. Whether one can scale to 10^5 points is strongly tied to the number of gridpoints per processor. For the problem class under consideration, 1000 to 10,000 points per processor is sufficient, given current-day parameters. Analysis indicates that all-to-all-based schemes are not so bad. Also, partitioning a domain into high-order quadrilateral (or hexahedral) elements can make for good local performance. A general-purpose gather-scatter code has been developed that supports 64-bit index sets. A coarse-grid solver has been developed that communicates with hundreds of members (not just the nearest members) and that speeds up processes by factors from 5 to 50 and makes an efficiency of 50% reasonable for a nodes-to-processors ratio of 7000.

In thermal hydraulics (TH), the turbulent entrance length has been established; LES-RANS equivalence has been established for cross-flow velocity distributions; and the recent success of a 37-pin analysis indicates that a design configuration of 217 pins is within reach.. LES and validated RANS are now being used to provide base velocity inputs to high-fidelity subcore coupled neutronics/TH simulations. A detailed analysis of 19- and 37-pin data are expected in January 2009. Simulation/analysis of a 217-pin case and detailed comparison to a reference experiment need to be conducted. Coupled TH/neutronics with detailed flow distributions need to be investigated with a whole-core model. And a core-scale upper plenum analysis of thermal striping phenomena needs to be conducted because boundary conditions have such a profound influence.

In computational science, a flexible and lightweight gs() communication utility is now enabling petascale deployment of many codes, and a new algebraic multigrid (AMG) coarse-grid solver has overcome a major impediment to scaling beyond 10,000 points. However, computational science is still going to be a major problem. Metadata and in situ running of VisIt are promising avenues to resolving this serious bottleneck. Otherwise, a ton of hardware is going to be needed. The clearest path to parallel memory access is through the distributed memory model. To date, a straight message-passing interface (MPI) is often the most efficient path to multicore usage. A radical change to programming model is only warranted through transformational paradigm shifts.

Berger asked what the next step was. Fischer replied that the million-cpu barrier would be crossed this year. There are many possible approaches, such as meshes from pin numbers. The users did not want a general-purpose mesh generator, just a single mesh. The CUBIT geometry and mesh generation toolkit was not designed for this type of problem.

Chen asked if single precision was good enough. Fischer replied, no, but double precision was coming out in the current month. Larger cache was needed. Failing that, one could put in only the active data.

Chalk asked if these activities were funded by the Laboratory-Directed Research and Development Program. Fischer replied that they were.

Alexandra Landsberg was asked to present an update on the Applied Math Program.

Applied Mathematics conducts research on mathematical models, methods, and algorithms to enable scientists to accurately understand complex physical, chemical, biological and engineered systems. It currently supports research activities in advanced linear algebra; discretization and meshing; multiscale, multiphysics systems; uncertainty quantification and error analysis; optimization; other research; and fellowships and workshops.

Approximately 75% of the FY08 $23.6 million goes to the national laboratories, and 25% to universities. The budget is growing, funding traditional areas and growing new areas.

Today’s program has a number of pending awards for multiscale mathematics and optimization for complex systems. Multiscale Mathematics has diverse temporal/spatial scales, multiple physical models; possibly many components (possibly dissimilar), complex connectivity (usually nonlinear), hard-to-predict behavior (often highly sensitive); complex systems analysis (e.g., combustion, materials, fluids, plasmas/MHD, and porous media). Optimization of Complex Systems is being grown in analysis and algorithms for stochastic optimization; theory and algorithms for very large, structured optimization problems; and analysis and algorithms for optimization problems with mixed variable types, including continuous, discrete, and categorical variables. Three panel reviews were convened in June 2008 that drew upon comments and advice from workshops. Two awards were made in 2008:

- Large-Scale Optimization for Bayesian Inference in Complex Systems and
- Next-Generation Solvers for Mixed-Integer Nonlinear Programs: Structure, Search, and Implementation.

Nineteen additional proposals are still "under consideration." Decision factors were based on peer-review comments, clearly identified new mathematical methods or algorithms, relevance to the DOE mission, and balance of breadth and depth.

The FY09 request of $43.2 million is 6.3 million over the FY08 appropriation. It will support a new joint Applied Mathematics–Computer Science Institute, a new effort in the mathematics of large datasets, and increases in key areas of long-term research.

Mathematics for the Analysis of Petascale Data was the subject of a workshop held June 3-5, 2008, in Rockville, Md. Innovative mathematical approaches and techniques are needed for finding the scientific knowledge in massive, complex datasets. The workshop goals were to understand the needs of various scientific domains, translate these needs into mathematical approaches and techniques, assess the current state-of-the-art, and target gaps and shortfalls that must be addressed. This topic resonates with activities at NSF and DARPA.

Setting up Joint Mathematics and Computer Science Institute(s) is also being looked at to focus on the challenges of computing at extreme scales that blur the boundaries between these disciplines. A workshop was held October 7-9, 2008, in Chicago with 32 participants. Key topics discussed were both computer science and mathematics oriented, and nine topics were selected as key problems:

- Inability to efficiently develop straight-forward, high-performance portable code;
- Using machines efficiently;
- Fault detection, tolerance, and management ;
- Sensitivities, uncertainty quantification, etc.;
- Multiscale/multiphysics;
- Fast implicit solves;
- Numerical stability of transient problems at scale;
- The untenability of debugging of correctness and performance; and
- Suboptimal algorithms for computer system resource management.

The workshop also identified cross-cutting tools, including a portable programming model; an execution model; algorithms for implicit methods, reformulations for larger time steps, and discrete optimization for page mapping, router management, etc.; mixed precision and reduced data representations; and libraries as test-beds/proofs-of-concept for new programming/execution models.

A joint Math/Computer Science Institute would have staffing at national laboratories and universities with a continuum of skills; approximately 10 to 20 members with a single principal investigator; a single theme with multiple projects; and an integrated mathematics and computer science effort. Funding of $1 million is too small; $3 million seems about right. Proposed work must clearly demonstrate the need for combined mathematics and computer science research to succeed. Milestones must depend on a joint effort.

New scientific capabilities that are too high risk to be carried out as business as usual need to be developed. A case-study approach will be used, focusing on fusion and accelerator physics, climate, combustion, and nuclear energy. High-risk/high-payoff projects will need to be defined. Success could provide a large increment in scientific capability. Types of risk include the well-characterized application of a new technology; well-established techniques applied to a new problem area; or fundamentally new approaches, particularly in domains where there is little prior art in modeling. For the four applications (fusion, climate, combustion, and nuclear energy), the timeline for progress over 3 to 5 years, the optimal end state in 10 years, the level of effort required to meet these goals, the organizational structure of a collaboration team, and external dependencies (e.g., SciDAC support) would need to be considered. Some cross-cutting issues have been identified: the need for petascale data infrastructure; institutes; robust and fast parallel I/O; program language support and kernel library support for multicore/ nonuniform memory access (NUMA) nodes; rapid prototyping tools; load balancing for new, large machines; and the continuation of SciDAC.

The Applied Mathematics PI Meeting was held Oct 15-17, 2008, at Argonne National Laboratory. More than 140 researchers attended. There were plenary talks on climate change, multiscale modeling, large-scale data analysis, towards exascale computing, uncertainty quantification and optimization, and advancing energy through algorithms. There were theme areas on linear and nonlinear systems, multiscale phenomena, uncertainty quantification/sensitivity analysis, and optimization of complex systems.

FY10 and beyond will be informed by numerous workshops and panel reports:

- Applied Mathematics at the U.S. Department of Energy: Past, Present, and a View to the Future
- Modeling and Simulation at the Exascale for Energy and the Environment Town Hall Meetings Report
- Mathematics for Analysis of Petascale Data Workshop Report
- Report on the Mathematical Research Challenges in Optimization of Complex Systems
- Multiscale Mathematics Initiative: A Roadmap

Program managers are actively working to define new research opportunities that build on and advance traditional ASCR Applied Mathematics strengths.

Stevens asked about the granularity of funding in mathematics and whether the average grant size were correct. Landsberg answered that the average size is $500,000, which is probably not the optimal size. The size distribution needs to be studied. That is one of the analytical axes that need to be used to study the program.

Omar Ghattas was asked to discuss uncertainty quantifications.

Key computational kernels for uncertainty quantification include: (1) estimate model parameters and their uncertainty from data (the statistical inverse problem) and (2) propagate parameter uncertainty through the model to predict quantities of interest and their uncertainty (the forward uncertainty propagation problem). The challenge is that the framework is often intractable for high-dimensional input parameter spaces and expensive forward problems, as is common for many problems in computational science and engineering described by partial differential equations (PDEs). The thesis of the presentation was that structure-exploiting methods, as commonly employed in deterministic inverse and PDE-constrained optimization problems, can help overcome the curse of dimensionality for the statistical inverse problem.

State estimation for atmospheric transport was introduced as a model problem, to motivate lessons learned from deterministic inverse problems that have the potential to inform and make tractable statistical inversion. The problem is ill-posed, so a regularization term has to be introduced. The state equation is posed in space and time, so the discretized optimality conditions produce a 100 billion by 100 billion system; to make this tractable, the state and adjoint equations are eliminated to produce a system in just the initial conditions. The question is: With measurements taken over a 3-hour window after an airborne release, can one reconstruct the initial conditions? The answer is, yes, fairly well.

Even though the system being solved is of dimension 100 million by 100 million the effective dimension is much smaller because the measurements provide information on just a handful of initial-condition modes. The conjugate gradient (CG) method is very effective for Hessians with such so-called "compact + identity" structure because these algorithms can tease out the low-dimensional information in a handful of iterations. Thus, inverse problems with 100 million parameters can be solved with as few as 20 iterations.

The Bayesian formulation for statistical inversion takes a forward model relating model parameters with observables and an associated uncertainty, actual observations and their uncertainty, and a prior estimate of model parameters and its uncertainty and seeks a statistical characterization of model parameters consistent with the observations, forward model, and prior model. The Bayesian formulation produces a so-called posterior probability density function (PDF) of the uncertain model parameters, which is easy to manipulate in one dimension but completely impenetrable in 100 (let alone 100 million) dimensions..

In the linear case and for Gaussian measurement noise, the posterior PDF for the model parameters is also Gaussian, and therefore its mean can be found by maximizing PDF, which is equivalent to solving the weighted-least-squares optimization problem. C_M^{-1} plays the role of the regularizer, and the posterior parameter covariance is given by the inverse of the Hessian. The "effectively-low-rank" structure of the Hessian permits fast (i.e., at a cost of a constant multiple of the forward-solution cost) estimation of the covariance and hence uncertainty in the model parameters.

In the nonlinear case, the posterior PDF is no longer Gaussian, and one is forced to sample the posterior using Markov-chain Monte Carlo (MCMC) methods. The resulting sample points can then be used to compute sample statistics, such as mean, covariance, etc. In common use, this approach breaks down when the model is even just moderately expensive, such as a 2D PDE or when the number of parameters is large. Even a few dozen parameters proves challenging. However, despite the non-Gaussianity, one would still expect that exploiting gradient and Hessian information would provide useful information, given the connection, at least locally, between the inverse Hessian and the parameter covariance matrix.

The Langevin Equation uses derivative information of the probability density to speed up sampling. (The sampling is carried out according to the Metropolis-Hastings criterion.) The inverse of the local Hessian can be introduced as a preconditioner for the Langevin equation to dramatically speed up sampling, resulting in the stochastic equivalent of Newton's method. When the posterior PDF is Gaussian, the stochastic Newton produces samples that act like independent draws from the true PDF, and thus *every* proposed move is accepted.

Some typical results were presented for 1-D inverse seismic wave propagation, where the problem is to reconstruct the shear modulus by discretizing the medium into a number of layers and then reconstructing the shear modulus of all of the layers from surface measurements of the reflected waveforms with a Bayesian-inversion framework. One result that emerged was that, for 65 parameters, the stochastic Newton method converges in several orders of magnitude faster time than a standard Metropolis Hastings MCMC method.

Experience has shown that, for linear statistical inverse problems, fast methods can be constructed that exploit the low-rank approximations of the Hessian. Hessian-preconditioned Langevin MCMC (i.e., stochastic Newton) is motivated by the connection to the deterministic Newton method; it exactly samples a Gaussian posterior; and it exhibits several orders of magnitude improvement over standard MCMC, even with a naive implementation. One can capitalize on several decades of advances in deterministic PDE-based optimization and inverse methods to vastly improve stochastic Newton by using inexact Newton ideas, trust-region methods, and/or exploiting the "compact + differential" structure of Hessians. The conclusion is that exploiting the deterministic PDE inverse problem structure is mandatory for scaling sampling methods like MCMC to high dimensions and expensive forward problems.

Collis asked about the structure of the Hessian of acoustic inverse problems. Ghattas replied that such problems often exhibit structure similar to the diffusion problem but for different reasons (in this case due to the band-limited nature of the observations).

Giles opened the floor to additional public comment. There being none, the meeting was adjourned for the day at 5:03 p.m.

Wednesday, October 29, 2008
Morning Session

The meeting was called to order at 8:37 a.m., and **Walter Polansky** was asked for an update on the SciDAC Program.

The SciDAC-2 goals are to create comprehensive, scientific computing software infrastructure to enable scientific discovery in the physical, biological, and environmental sciences at the petascale and to develop a new generation of data management and knowledge discovery tools for large data sets (obtained from scientific users and simulations).

A 2003 memorandum states Dr. Orbach's intentions regarding management of SC SciDAC activities in each of the program budgets. The policy states that the SciDAC Director is "to review and sign-off on SciDAC FWPs, grant initiations and renewals, and AFP changes and program guidance letters concerning formulation and execution of budgets." Michael Strayer is the SciDAC Director.

Some awards were made at the end of FY06, and a long continuing resolution delayed funding for FY07, postponing the startup of several centers and institutes. In FY07, a new RFP was issued for accelerator science and simulation, funded by ASCR, SC, NNSA, and NSF. The Centers for Enabling Technology, the Outreach Center, the Scientific Applications Partnerships, and the institutes were 100% funded by ASCR. The hardware component of SciDAC was never funded.

The upcoming review of SciDAC will take an ensemble perspective. The science applications will be reviewed from January to April of 2009, and the centers and institutes will be reviewed from April to May. The process will use peer review panels, and the reviewers will reflect the interests of the funding agencies and offices. The intention is to deliver a recommendation package for the entire program to the Program Director.

The SciDAC Coordination Group is comprised of program managers from ASCR, BER, Basic Energy Sciences, Fusion Energy Sciences, High Energy Physics, Nuclear Physics, and NNSA. The Coordination Group has formulated the notice and announcement, managed peer reviews, recommended projects for funding, prepared award packages, managed awarded projects, and developed procedures for organizing, scheduling, and conducting peer reviews of the SciDAC portfolio.

Bailey asked about the correlation between SciDAC awards and allocations of computer time. Polansky said that there is a level playing field. SciDAC projects have to compete for allocations but are more computationally ready because of the SciDAC institutes and centers. The NERSC allocation process is different from INCITE and helps DOE projects meet their missions. However, NNSA work generally cannot be put on NERSC.

Negele asked if the relationship between SciDAC and INCITE would be discussed in the upcoming review of SciDAC projects. Polansky replied that INCITE complements SciDAC and vice versa, so the relationship will not be the primary purpose of the review. However, the relationship is likely to come up on specific SciDAC projects. Some SciDAC applications rely on INCITE to provide key resources needed to advance science through computation.

Simon stated that, in the future, it would be helpful for NERSC to automatically grant allocations to projects that have been selected by the various INCITE and SciDAC competitive reviews. Berger noted that the concern about the reviews has been that there is double jeopardy. Polansky noted that, if people want access to the LCFs, they have to go through the INCITE process.

Negele noted that the application and review processes take important time away from science. Streamlining the processes would free up significant time for the review panelists. Polansky replied that there will be opportunities in this review to streamline the portfolio. However, there are no funds to seize those opportunities.

Chen observed that there may be more readiness variability during the conduct of a project because code will have to be developed during the project.

Simon asked what SciDAC X will look like. Polansky replied that the Mathematics–Computing Science Institute may be one component. SciDAC X may be a more fluid process; it may not be driven by 14 science projects but by 5 centers.

White asked how well the Program was doing on metrics. Polansky replied that there is strong evidence that other nations are trying to copy SciDAC. Many prominently cited computational breakthroughs have a SciDAC relationship. How SciDAC advances science needs to be explained to a broader audience.

Lali Chatterjee was asked to report on the exascale workshops, which are bringing together science communities from a range of DOE offices to explore the challenges that can be addressed with exascale computing. A series of DOE-sponsored exascale workshops will typically have opening and plenary sessions, break-out sessions, interconnectivity/overview groups, writing-team meetings, and a workshop report that identifies the science community's needs for exascale computing. The workshops will address the key questions: What are the science grand challenges? Why is exascale computing needed to help solve them? And what are the priorities? The workshops will build on prior studies that have identified important challenges across science.

The first workshop, on Nov. 6–7, 2008, will deliberate the challenges in climate change science and the role of computing at the extreme scale. A 2008 BER Workshop identified three grand challenges:

- Characterizing the Earth's current climate and its evolution over the past century to its present state,
- Predicting regional climate change for the next several decades, and
- Simulating Earth system changes and their consequences over centuries.

This community is poised to step up to the exascale in such areas as model development and integrated assessment; algorithms and the computational environment; data, visualization, and productivity; and decadal predictability and prediction.

Other workshops in an advanced stage of development are those on high energy physics (Dec. 9–11, 2008, SLAC National Accelerator Laboratory) and nuclear physics (Jan. 26–28, 2009, Washington, D.C.). These will address the fundamental science questions identified in the report, *Connecting Quarks with the Cosmos.*

More 2009 workshops are being planned on nuclear energy, fusion science, biology, and material science and chemistry. Addressing such topics usually requires a holistic approach. For example, the physics workshops will consider theory, experiment, simulation, astrophysics, accelerator science, cosmology, ESnet, and high-performance computing.

These exascale workshops will help define why exascale computing is needed and where, when, and how exascale computational science will be part of the synergy that will solve the outstanding science grand challenges of tomorrow.

Giles asked what was expected to be learned from the workshops that will affect exascale computing. Chatterjee replied that people are being brought in who are closely linked to exascale computing and who will influence high-performance computing. Stevens commented that the question is whether a strong-enough science case can be constructed to justify the construction of hardware and the development of software. Chatterjee added that each community is different in needs and likely workshop outcomes.

Manteuffel pointed out that the fusion program has some components that can be addressed by computing and some that cannot. Exascale alone will not solve the problem. The parts that need exascale computing and advances in computer science and mathematics need to be identified. Giles stated that these workshops should reveal these needs together. Chatterjee said that that was part of the holistic approach. Helland added that people from mathematics and computer science will be part of these workshops and that there will be follow-up workshops on computer science and mathematics.

White asked what was being focused on. Chatterjee said that the focus will be on identifying science challenges that really need exascale computing, but there will be ancillary benefits. Stevens added that, if one looks at the disciplines, one has to ask, "What are the key questions that can be addressed?" and then, among the tractable questions, "Which can only be addressed with exascale computing?"

Michael Strayer was asked to respond to the INCITE Committee of Visitors (COV) recommendations.

The COV that assessed the INCITE Program recommended that the selection processes for leadership-class and DOE capability-class computing should be separated and that a significant portion of INCITE computational resources be allocated to high-end DOE capability-class computing through an INCITE-type process. This recommendation seems to have been inspired by the recent shortfall of resources at NERSC. The Office intends to increase the pace of upgrades at NERSC and to provide even more computational resources to SC computational efforts by allocating older, but still scientifically useful, leadership resources through the NERSC allocation process. The NERSC 6 upgrade has just been completed, and planning for NERSC 7 will bring the facility to near-LCF computing.

The Office concurs with the COV that INCITE awards should be fewer in number and larger in size with the expectation of demonstrated concurrency across a very large number of cores and will work to re-balance the INCITE portfolio. Projects deemed important but not ready are referred to the SciDAC Outreach Center and may be given access for scaling efforts from the facility director reserve to bring computations up to speed.

The Office agrees that renewal should meet an achievement threshold and will develop criteria for renewals that will be implemented with the calendar year 2009 INCITE allocations. These criteria will be broadly communicated to the community.

The Office concurs with the COV that INCITE should continue to provide robust expert assistance to the science teams performing leadership-class computing, and a good deal of the discussions with LCF management has focused on user assistance.

The Office agrees that the selection process should be made as transparent and as uniform across disciplines as is practical and will put the selection-criteria information to implement this recommendation on the next year's INCITE proposal website and will further stagger the INCITE and NERSC annual calls. As a result, the 2010 INCITE Call for Proposals will open in mid-April 2009 and close in mid-July 2009 to relieve the burden on submitters and reviewers.

In response to the recommendation that the computational readiness review process should adopt a more descriptive outcome and a more systematic process the Office is implementing an overall grading scheme for the computational-readiness review with a scale from not ready (1) to ready (5). This scheme will not solve all of the problems, however. The computational readiness review already includes a panel review composed of computational experts from DOE's LCFs and NERSC. The option of adding outside computational experts to this panel will be explored.

As recommended by the COV, an appeals process for allocation decisions will be developed and implemented to supplement the current appeals process, which is approved by the Office of Management and Budget (OMB). The SC Director's reserve (10% of the facilities) will be used for this process.

This COV was very effective, and its help was very welcomed by the Office.

Strayer issued a new charge to the Committee for a COV for the Computer Science group. The COV is being asked, for both the DOE laboratory projects and the university projects, to assess the efficacy and quality of the processes used (1) to solicit, review, recommend, and document proposal actions and (2) to monitor projects' and programs' activities. The COV is also being asked to comment on how the award process has affected (1) the breadth and depth of portfolio elements and (2) the national and international standing of the program with regard to other computer science research programs focused on high-performance scientific computing and analysis of petascale datasets. A draft report will be due to ASCAC at its August 2009 meeting.

Berger thanked the Office for its help to the COV. INCITE is a remarkable program and has been stood up rapidly. She asked if there could be an INCITE-like allocation process for SciDAC projects and whether the office could collect more information to make the case for exascale computing. Strayer replied that the Office is currently evaluating how to eliminate the double-jeopardy situation. Currently, the issue is not fully understood. It seems like the second review is just checking off a box. Berger pointed out that the first review does not provide the successful SciDAC projects the resources they need to carry out their projects. Strayer noted that these resources are incredibly valuable. It may be that a SciDAC project needs these resources. It may be possible to work this assessment into the SciDAC review. SciDAC *does* require a need for advanced computational science. INCITE and SciDAC should be interchangeable.

Bailey pointed out that the two programs do not approach this computational need in the same way. If the SciDAC review were to satisfy the INCITE review, the assessment of computer science need and readiness would have to be identical.

Negele stated that some of the redundancy could be eliminated and that the process could be made more efficient. Strayer pointed out that the SciDAC COV stated that there were inequities but that the outcomes were excellent. At the time, the resources were not in hand and then were never funded, producing stress on NERSC. At this point, the Office is trying to recover from this history and trying to get resources into the SciDAC Program. It is expected that the reviews will reveal the needs of the projects.

As a public comment, Hendrickson pointed out that the results of the SciDAC review were announced just weeks before the deadline for INCITE proposals. His project, the Combinatorial Scientific Computing and Petascale Simulations (CSCAPES) Institute, had to scramble for resources. The situation is complicated by the facts that the vagaries of panels will lead to different outcomes and that there are different metrics for the two reviews.

A break was declared at 10:18 a.m. The meeting was called back into session at 10:38 a.m. to consider the new charge to the Committee (a copy of which had been distributed to the Committee) and to continue the discussion of the balance report.

Berger asked what the focus of the second point of the charge was. She wanted to know what part of the high-performance-computing portfolio was being left out. Chalk said that that statement was intended to ensure that the ASCR program was compared with similar programs. It was not intended to exclude any program elements. Strayer added that a letter could go to the Chair that lists all the elements to be considered.

Giles asked who should be considered for membership on this COV. Simon suggested that those responsible for large parts of this portfolio should not be part of this COV. The chair should be an academic. Strayer suggested that recommendations be e-mailed to Chalk and Giles. Manteuffel suggested that the chair be an ASCAC member to lend experience, focus, and communication. Strayer pointed out that conflicts of interest must also be considered. Because people who are not Committee members would not be special government employees, they would not have as many restrictions. Chalk said that the selected members will be announced on the Committee's website

Giles proposed the formation of a new subcommittee to finalize ASCAC's response to the charge with a report due at the August 2009 meeting and an interim status report at the March 2009 meeting. That subcommittee might consider the interpretation of the charge, building on the existing subcommittee report as well as new inputs, and weighing key issues raised in recent discussions (e.g., balance in the context of strategic plans, leadership, decision-support models, and support for applications). The subcommittee would not need to answer all these questions but should identify critical issues.

Bailey asked whether the balance needs to be changed to accommodate a strategy or whether the balance drives the implementation of the strategy. Berger said that how the subcommittee proceeds should be informed by discussions with the Director of the Office. Strayer noted that the discussions of the previous day touched on the dimensions of the issues involved and how they are intertwined. Berger suggested that the Committee Chair and the Director should set the scope and subject of the subcommittee's deliberations.

White asked what was meant by "building on the existing subcommittee report" and whether that meant building on the accumulated information and data. Giles replied, yes.

Manteuffel said that the Chair and Director should be careful in writing the charge. The topic is dynamic, and the assessment has to be able to change from time to time, also.

White asked if the charge could be changed. Giles reflected that, in the past, there have been charges that had faded away. It is not desirable to do that anymore. It is possible to have a report that says that the topic was too large, and here is what was learned about the subtopics that could be addressed. There should be some continuity between the previous subcommittee and the new subcommittee. Simon suggested that the new chair should be someone who sat on the original subcommittee.

Giles asked if there were conflict-of-interest issues to be addressed. Negele said that coincidence of interest should be focused on rather than conflict of interest.

Bailey asked how much of the outyear budgets the subcommittee could see. Strayer stated that that information is embargoed. Chalk said that the subcommittee meetings are not public, and the Office will provide as much information as possible. Negele said that addressing budget scenarios can be done without embargoed information. Stevens noted that analysis is easier if there are strategies to be resolved for the different scenarios. The subcommittee needs to take a shot at a conceptual analysis, bringing in the science to drive the logic.

Bruce Hendrickson was asked to discuss graphs and high-performance computing.

The term computational science and engineering brings to mind partial differential equations and numerical methods. But combinatorial algorithms have long played a key enabling role in sparse direct methods and preconditioning, load balancing and architecture exploitation, optimization and uncertainty quantification, and mesh generation. Graphs feature strongly in emerging application areas, such as biological networks, chemistry, and advanced data analysis.

Some examples in sparse-matrix methods include re-orderings for sparse solvers, which are powerfully phrased as graph problems; graph partitioning, graph traversals, and graph eigenvectors; data structures for efficient exploitation of sparseness; derivative computations for optimization (e.g., matroids, graph colorings, and spanning trees); and preconditioning for iterative methods, which include incomplete factorizations and graph partitioning for domain decomposition. Graph coloring is a different problem; it identifies a set of vertices that are nonrelated and is used in parallel adaptive meshing and automatic differentiation.

Parallelizing scientific computations is aided by graph algorithms, geometric algorithms, and reordering for memory locality.

The Combinatorial Scientific Computing community was formed to develop, apply, and analyze combinatorial algorithms that enable scientific and engineering computations. Four international workshops have been held, and special issues of journals have been published. The SciDAC CSCAPES Institute was established to study combinatorial scientific computing and petascale simulations.

Architectural challenges for graphs include the runtime being dominated by latency, essentially no computation to hide memory costs, and an access pattern that is data dependent, producing potentially abysmal locality at all levels of the memory hierarchy.

Desirable architectural features would be low latency and high bandwidth, especially for small messages; latency tolerance; light-weight synchronization mechanisms; and a global address space. There is one machine with these properties, the Cray MTA-2 and its successor, the XMT.

The MTA-2 has a cacheless architecture. It achieves latency tolerance via massive multi-threading: Each processor has hardware support for 128 threads with a context switch in a single tick, global address space (hashed to reduce hotspots), and multiple outstanding loads. However, the MTA-2 is old, its clock rate is 220 MHz, and the largest machine is 40 processors.

The problem of finding the shortest path between two vertices is a standard graph operation. A parallel algorithm for this was implemented in MPI and run on the IBM/LLNL BlueGene/L with a 4-billion-vertex, 20-billion-edge Erdös-Renyi random graph. It ran in 1.5 seconds on 32,000 processors. A simple analysis shows that this algorithm need only touch about 200,000 vertices. A similar problem was run on an MTA-2 with 32 million vertices and 128 million edges and took 0.7 second on one processor and 0.09 second on 10 processors. The algorithm touched about 23,000 vertices. Thus, the one-processor MTA-2 run took about half as long as the BG/L run, while visiting one-eighth as many vertices. This result suggests that, given enough memory, a handful of MTA-2 processor can solve this problem as fast as 32,000 BlueGene/L processors. One can quibble about the details of this comparison, but it highlights the dramatic difference that novel architectures can have for some classes of computations.

The MTA-2 suggests an alternative model for multi-core-node programming with shared memory with a simple programming model on the nodes, latency tolerance, and fine-grained parallelism and dynamic, fine-grained load balancing. There are still many open questions at the interface of math and computer science: How best to build and program multiple cores? Is there a unified programming model that achieves high inter- and intra-node performance? And how to get from here to there? Graph algorithms can serve as a canary in a coal mine for new architectures, languages, and programming environments that would allow stressing systems in ways that anticipate the needs of future applications

Existing HPC applications are getting more complex in terms of unstructured and adaptive grids, multiscale and multiphysics, and complex data structures and dependencies. Emerging applications (e.g., data analysis, biological networks, and decision support) are even more demanding. But memory performance is not increasing as fast as processors are. As a result, latency will be increasingly important.

Simon asked if the group had any experience with the XMT. Hendrickson replied that the group has access to a machine at Cray. Much of the MTA-2 was expensive custom hardware. The XMT was designed to use a lot of hardware components that were already developed, and so is much more affordable. Initial results show good scalability for problems with a high degree of thread-based parallelism.

Stevens asked if there were a correlation between graphs and the GUPS [giga-updates per second] benchmark. Hendrickson replied that there was. However, the way the GUPS benchmark is structured allows for implementations that do not stress the network latency as much as the benchmark developers intended. As a result, Blue Gene and Red Storm do well on GUPS. Stevens asked if the group had thought about developing a new benchmark that could not be cheated on. Hendrickson replied that some have tried to develop graph-theoretic benchmarks

The floor was opened to public comments. There being none, the meeting was adjourned at 11:46 a.m.

Respectfully submitted,
Frederick M. O'Hara, Jr.
Recording Secretary
November 17, 2008